

Constraining Pulsar Magnetospheric Electrodynamics via Joint Radio and γ -ray LC Fitting

¹Seyffert A.S., ¹Venter C., ²Harding A.K., ^{2,3,4}Kalapotharakos C.,
on behalf of the *Fermi* Collaboration

¹Centre for Space Research, North-West University, Private Bag X6001, Potchefstroom 2520, South Africa

²Astrophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

³Universities Space Research Association (USRA), Columbia, MD 21046, USA

⁴University of Maryland, College Park (UMDCP/CRESST), College Park, MD 20742, USA

albertus.seyffert@gmail.com, christo.venter7@gmail.com

Abstract

The wealth of γ -ray pulsar data the *Fermi* Large Area Telescope has delivered has stimulated the development of sophisticated magnetospheric and radiation models that promise powerful dual-band constraints on a number of pulsar properties. Investigations using statistical light curve fitting have, however, been hampered by the large disparity between the radio- and γ -ray-band errors. We introduce the Scaled-Flux Normalised χ^2 test to overcome this, and use it to search for the correlation between magnetospheric conductivity and pulsar spin-down luminosity reported by Brambilla et al. (2015).

Introduction

To date the *Fermi* Large Area Telescope (LAT) has reported the detection of pulsed emission from at least 216 γ -ray pulsars, many of which also pulsate in the radio band. This wealth of dual-band data has stimulated the development of sophisticated models of the pulsar magnetosphere and the mechanisms responsible for the pulsed emission. Crucially, many of these models yield predicted light curves (LCs) in both bands concurrently, assuming a single underlying magnetic field. These model LCs can, in principle, be used to constrain, among other things, the viewing geometries of the observed *Fermi* LAT pulsars via statistical goodness-of-fit testing.

Doing so using the traditional χ^2 test statistic, however, results in constraints (and best-fit LCs) that are dominated by the radio-band fit. This is caused by the large disparity between the radio and γ -ray LC errors, which essentially implies a weighting for the single-band χ^2 values that leans heavily in favour of the radio fit. The coloured bars in the left margin are a to-scale representation of this weighting for PSR J0631+1036 (our exemplar for this poster). How this weighting arises, and what the colours of the bars mean, are explained in the sections that follow. Figure 2 (c) shows the resulting radio-dominated concurrent fit.

Statistical fitting

In the **single-band** case, the traditional χ^2 goodness-of-fit test (as applied to LC fitting) minimises the test statistic

$$\chi^2(\mathbf{M}) = \sum_i \left(\frac{D_i - M_i}{\varepsilon_i} \right)^2$$

over the predicted LCs \mathbf{M} a given model produces as possible fits to an observed LC, \mathbf{D} , with errors ε . The subscript i denotes a value for the i th phase bin, and the summation is done over all n_{bins} phase bins in the observed LC (see Figure 3 for example LCs).

For a model with n_p free parameters the χ^2 value of a good fit is χ_n^2 distributed, where $n = n_{\text{bins}} - n_p$ is the number of degrees-of-freedom of the fit, and has an expected value of n . For a discussion on how this distribution translates into constraints on the model parameters (like those shown in Figure 2), see Avni (1976).

In the **dual-band** case, where a model produces concurrent radio and γ -ray LCs, the combined χ^2 value is $\chi_c^2 = \chi_r^2 + \chi_\gamma^2$.

Radio dominance

Since the LC errors appear in the denominator in the definition of χ^2 , they determine the sensitivity of the single-band χ^2 value to deviations from a good fit. If the relative errors in the two bands are too disparate, the dual-band fit will be dominated by the χ^2 value of the more sensitive band. The coloured bars in the left margin demonstrate this sensitivity for the two bands for PSR J0631+1036. For both bars black indicates a good fit, while white indicates no fit (using the background level \mathbf{B} as the model LC).

Attempted solutions

Johnson et al. (2014) addressed the radio dominance by scaling the radio LC errors for each concurrent fit using the average relative error for the on-peak bins of the corresponding γ -ray LC. For some pulsars, however, this scaling was not enough, and they needed to adjust the scaling factor by hand.

Our new test statistic

To resolve the radio dominance issue it is necessary to reframe the problem of statistical fitting as a geometric one.

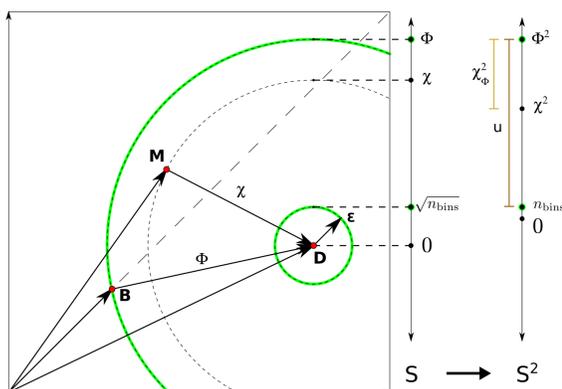


Figure 1: The LCs on the outer green sphere in n_{bins} -dimensional LC space fit \mathbf{D} just as well as the background level \mathbf{B} ($\chi_\phi^2 = 0$). The LCs on the inner green sphere are all statistically good fits to \mathbf{D} ($\chi^2 = n$; $\chi_\phi^2 = 1$). These two fiducial spheres, which exist for both bands, allow the definition of an unambiguous shared unit in χ^2 space, u . A fit's χ_ϕ^2 value is a measure of the distance between \mathbf{M} and \mathbf{D} using this unit.

The geometry of the χ^2 test

An LC can be thought of as a point in n_{bins} -dimensional LC space, with the intensity in each phase bin being its coordinate along one of those dimensions (see Figure 1). If the errors on the observed LC are used

as a basis for this space, they essentially define the units in which distances are measured. In this context, the χ^2 test statistic simply measures the Euclidean distance between two LCs in units of the LC errors. The single-band χ^2 values, therefore, each correspond to a distance in a separate LC space, and carry with them implicit units which make it improper to add them together to obtain χ_c^2 .

The Scaled-Flux Normalised χ^2 test statistic

To compensate for these implicit units we define an explicit shared unit $u = \Phi^2 - n$ in χ^2 space, where $\Phi^2 = \chi^2(\mathbf{B})$ is the scaled flux, or the χ^2 value of the background \mathbf{B} (see Figure 1). We then express the goodness-of-fit as

$$\chi_\phi^2 = \frac{\Phi^2 - \chi^2}{u}$$

The expected value for a good fit is $\chi_\phi^2 = 1$, while for fits equivalent to the background, $\chi_\phi^2 = 0$. This normalisation exactly equalises the weights assigned to the radio and γ -ray fits. Coincidentally, this normalisation allows us to quantify the weights implied by the LC errors, and compare the combined values for the two test statistics:

$$\chi_{\phi,c}^2 = \frac{1}{2} \chi_{\phi,r}^2 + \frac{1}{2} \chi_{\phi,\gamma}^2, \quad \text{while} \quad -\chi_c^2 \sim \frac{u_r}{u_c} \chi_{\phi,r}^2 + \frac{u_\gamma}{u_c} \chi_{\phi,\gamma}^2,$$

with $u_c = u_r + u_\gamma$. Additionally, normalising the combined χ^2 value in this way is equivalent to rescaling the (uniform) radio LC error as

$$\varepsilon_R = \sqrt{\frac{\Phi_r^2}{u_r + n_r}} \varepsilon_r.$$

Important properties of χ_ϕ^2

- $\chi_\phi^2 = 1$ for a good fit.
- $\chi_\phi^2 = 0$ for a background-equivalent fit (a.k.a. a null fit).
- χ_ϕ^2 values obtained for different bands can be combined by averaging.
- Allows direct goodness-of-fit comparison across both pulsars and models, for both single-band and dual-band fits (respectively).
- Allows direct goodness-of-fit comparison across bands.
- Constraints on model parameters can be sensibly obtained even if a model's best fit is not a good fit.
- Retains all the important statistical properties of the traditional χ^2 test.
- In the limit where the radio and γ -ray LC errors are comparable, $\chi^2 \sim -\chi_\phi^2$.

An example fit

Figures 2 and 3 show the results of both single-band and dual-band fits for PSR J0631+1036. In Figure 2 we see that the traditional χ^2 -combined dual-band fit (panel (c)) is completely dominated by the radio-band fit (panel (a)), with the cross-diagonal structure of the γ -ray-only goodness-of-fit map (panel (b)) completely absent. The χ_ϕ^2 -combined fit (panel (d)), on the other hand, is sensitive to the structure of both single-band goodness-of-fit maps.

The best-fit LCs in Figure 3 demonstrate the effect of the increased sensitivity of the χ_ϕ^2 -combined fit to the γ -ray goodness-of-fit.

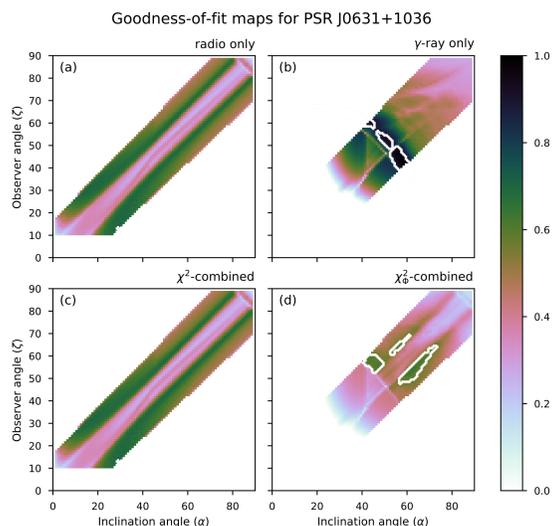


Figure 2: The goodness-of-fit maps for both single-band and dual-band fits for PSR J0631+1036. The χ^2 -combined map, (c), is scaled in an analogous fashion to the other maps to allow comparison. The colour scale is the same across all 4 panels. The white contours constitute Monte-Carlo-derived constraints on the pulsar's viewing geometry. Note the very small contour for the radio and χ^2 -combined maps (at $\sim (30, 10)$).

Applying the χ_ϕ^2 test to a real problem

As a first test of this new test statistic we investigate the correlation reported by Brambilla et al. (2015) between the conductivity, σ , of a pulsar's magnetosphere and its spin-down luminosity \dot{E} . They used the so-called Force-Free Electrodynamics Inside, Dissipative Outside

(or FIDO) model for the magnetosphere (Kalapotharakos et al., 2014), which assumes infinite conductivity inside the light cylinder, and finite conductivity outside.

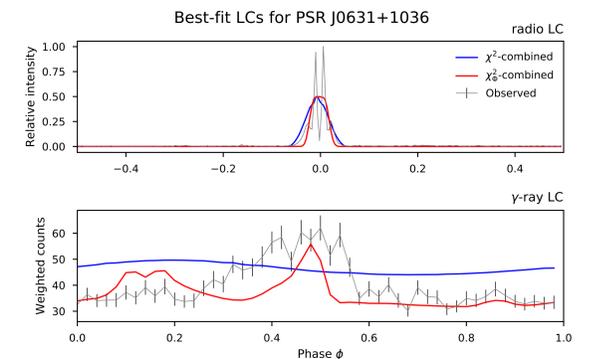


Figure 3: Radio and γ -ray best-fit LCs for dual-band fits to PSR J0631+1036. Notice the comparatively small errors on the radio LC.

They compared both LCs and curvature radiation spectra produced by this model to LCs and phase-resolved spectra observed by *Fermi* LAT for 8 particularly luminous pulsars. They found that increasing \dot{E} correlates with increasing best-fit σ .

As a crude dual-band model, we assume a hollow cone model for the radio emission, and the FIDO model for the γ -ray emission. We then find γ -ray-only and concurrent fits for 65 *Fermi* LAT pulsars (of which 34 are radio-loud), with 5 free parameters: σ , LC amplitude, and 3 parameters for the viewing geometry. The selected pulsars all have $\Phi^2/n_{\text{bins}} > 10$.

While the results (pictured in Figure 4) are still only preliminary, we find that the previously reported correlation is not present in either the γ -ray-only fits, or the concurrent fits. It should be noted, however, that this initial investigation was done over a parameter space that differs markedly from that explored by Brambilla et al. (2015). Specifically, their grid for σ is [0.3, 1, 3, 5, 10, 30, 60], while ours is [1, 10, 100, 1000, 10000].

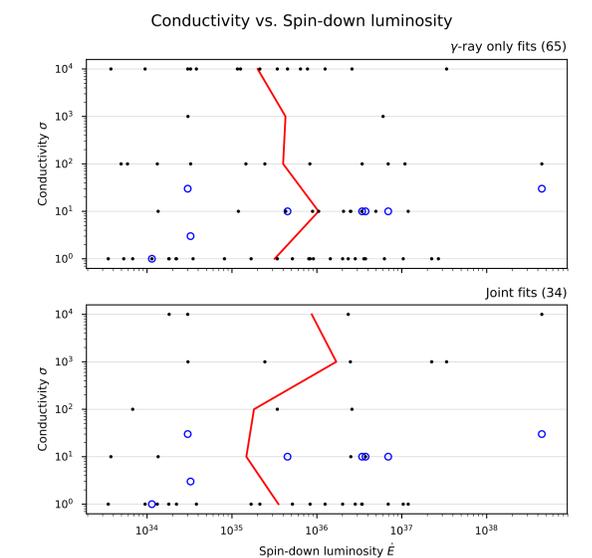


Figure 4: Scatter plots of best-fit σ vs. \dot{E} for 65 *Fermi* LAT pulsars (the black dots). Joint fitting was done for 34 of the pulsars (the radio-loud subset; *bottom*), and γ -ray-only fitting was done for all 65 pulsars (*top*). The red line in each plot connects the goodness-of-fit weighted averages of the black points at each value for σ in our grid. The blue circles are the fits found by Brambilla et al. (2015).

Conclusion

The Scaled-Flux Normalised χ^2 test (χ_ϕ^2) enables robust single-band, dual-band (and potentially multi-band) statistical LC fitting, and yields goodness-of-fit values that are easily comparable across pulsars and models, as well as across bands. The combined fits it yields are sensitive to the goodness-of-fit in both bands, even when a large LC error disparity exists.

Our investigation into the previously reported correlation between σ and \dot{E} is promising, but needs to be refined. Brambilla et al. (2015) used the FIDO model to fit curvature radiation spectra, which may be a more sensitive probe of σ , and their grid focused on lower σ cases. A good next step for us would be to change our σ grid, and incorporate spectral fits into our results.

References

- Avni, Y. 1976, ApJ, 210, 642
Brambilla, G., Kalapotharakos, C., Harding, A. K., & Kazanas, D. 2015, ApJ, 804, 84
Johnson, T. J., Venter, C., Harding, A. K., et al. 2014, ApJS, 213, 6
Kalapotharakos, C., Harding, A. K., & Kazanas, D. 2014, ApJ, 793, 97